

07-27-00
HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

**UTILITY
PATENT APPLICATION
TRANSMITTAL**
*(Only for new nonprovisional applications under
37 C.F.R. 1.53(b))*

Attorney Docket No.	2789.2006-001
First Named Inventor or Application Identifier	William J. Dally
Express Mail Label No.	EL290724694US

A
C880
07/26/00
C880
09/625802
U.S. PTO
07/26/00

Title of Invention

HIGH-SPEED, LOW-POWER CROSSBAR SWITCH

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. <input type="checkbox"/> Fee Transmittal Form <i>(Submit an original, and a duplicate for fee processing)</i>	6. <input type="checkbox"/> Microfiche Computer Program <i>(Appendix)</i>
2. <input checked="" type="checkbox"/> Specification [Total Pages [25]] <i>(preferred arrangement set forth below)</i>	7. <input type="checkbox"/> Nucleotide and/or Amino Acid Sequence Submission <i>(if applicable, all necessary)</i>
- Descriptive title of the invention	a. <input type="checkbox"/> Computer Readable Copy
- Cross References to Related Applications	b. <input type="checkbox"/> Paper Copy <i>(identical to computer copy)</i>
- Statement Regarding Fed sponsored R & D	[] Pages
- Reference to microfiche Appendix	c. <input type="checkbox"/> Statement verifying identity of above copies
- Background of the Invention	
- Summary of the Invention	
- Brief Description of the Drawings	
- Detailed Description	
- Claim(s)	
- Abstract of the Disclosure	
3. <input checked="" type="checkbox"/> Drawing(s) (35 U.S.C. 113) [Total Sheets [7]] [] Formal [X] Informal	8. <input type="checkbox"/> Assignment Papers <i>(cover sheet & documents)</i>
4. <input type="checkbox"/> Oath or Declaration/POA [Total Pages []]	9. <input type="checkbox"/> 37 C.F.R. 3.73(b) Statement <input type="checkbox"/> Power of Attorney <i>(when there is an assignee)</i>
a. <input type="checkbox"/> Newly executed <i>(original or copy)</i>	10. <input type="checkbox"/> English Translation Document <i>(if applicable)</i>
b. <input type="checkbox"/> Copy from a prior application (37 C.F.R. 1.63(d)) <i>(for continuation/divisional with Box 17 completed)</i>	11. <input checked="" type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 <input checked="" type="checkbox"/> Copies of IDS Citations
i. <input type="checkbox"/> <u>DELETION OF INVENTOR(S)</u> Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. 1.63(d)(2) and 1.33(b).	12. <input type="checkbox"/> Preliminary Amendment
5. <input type="checkbox"/> Incorporation By Reference <i>(useable if Box 4b is checked)</i> The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.	13. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) <i>(Should be specifically itemized)</i>
17. If a CONTINUING APPLICATION , check appropriate box and supply the requisite information: [] Continuation <input type="checkbox"/> Divisional <input type="checkbox"/> Continuation-in-part (CIP) of prior application No.: _____	14. <input type="checkbox"/> Small Entity <input type="checkbox"/> Statement filed in prior application, Statement(s) status still proper and desired
Prior application information: Examiner: _____	15. <input type="checkbox"/> Certified Copy of Priority Document(s) <i>(if foreign priority is claimed)</i>
	16. <input type="checkbox"/> Other: _____

ACCOMPANYING APPLICATION PARTS

8. Assignment Papers *(cover sheet & documents)*
9. 37 C.F.R. 3.73(b) Statement Power of Attorney
(when there is an assignee)
10. English Translation Document *(if applicable)*
11. Information Disclosure Statement (IDS)/PTO-1449 Copies of IDS Citations
12. Preliminary Amendment
13. Return Receipt Postcard (MPEP 503)
(Should be specifically itemized)
14. Small Entity Statement filed in prior application, Statement(s) status still proper and desired
15. Certified Copy of Priority Document(s)
(if foreign priority is claimed)
16. Other: _____

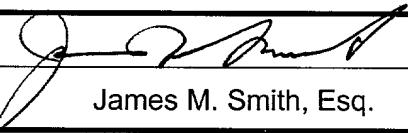
17. If a **CONTINUING APPLICATION**, check appropriate box and supply the requisite information:

[] Continuation Divisional Continuation-in-part (CIP) of prior application No.: _____

Prior application information: Examiner: _____ Group Art Unit: _____

18. CORRESPONDENCE ADDRESS

NAME	James M. Smith, Esq. HAMILTON, BROOK, SMITH & REYNOLDS, P.C.				
ADDRESS	Two Militia Drive				
CITY	Lexington	STATE	MA	ZIP CODE	02421-4799
COUNTRY	USA	TELEPHONE	(781) 861-6240	FAX	(781) 861-9540

Signature		Date	7/26/00
Submitted by Typed or Printed Name	James M. Smith, Esq.	Reg. Number	28,043

Date: 7-26-00 Express Mail Label No. EL 2907246 94 US

Inventor(s): William J. Dally and Daniel K. Hartman
Attorney's Docket No.: 2789.2006-001

HIGH-SPEED, LOW-POWER CROSSBAR SWITCH

RELATED APPLICATION(S)

This application claims the benefit of U.S. Provisional Application No. 60/198,188, filed April 19, 2000, the entire teachings of which are incorporated herein 5 by reference.

BACKGROUND OF THE INVENTION

At the heart of telecommunications switches and network routers is a switching fabric that is often constructed from one or more crossbar switches. Such switches forward packets between input ports and output ports in network routers and connect 10 circuits from input ports to output ports in telecommunications switches.

In the past, high-speed switches, over 1Gb/s per port of bandwidth, were typically implemented in bipolar or GaAs technology. While this fast circuit technology enabled construction of a fast crossbar switch, this technology is expensive and does not permit the integration of CMOS logic functions on the same chip with the 15 switch.

Most CMOS crossbar switches use static gate circuits to implement each crosspoint. Tri-state inverters may be employed to selectively gate inputs onto an output bus, or, alternatively, a multiplexer constructed from CMOS gates may be used at each output to select the active input. While these circuits are simple and robust, they 20 also require considerable chip area, dissipate large amounts of power, and have limited scalability. A conventional crossbar of this type is described, for example, in U.S. Patent 4,914,429.

An improved CMOS crossbar switch design was reported by Shin and Hodges (Shin and Hodges, *A 250-Mbit/s CMOS Crosspoint Switch*, JSSC 28(2), April 1989, pp. 478-486). Their switch employs a static CMOS inverter for a crosspoint but improves efficiency by reducing the swing of the column (output) line of the switch using a

5 CMOS inverter with resistive feedback. While an improvement, this switch requires large tri-state inverters at the crosspoints. This type of switch circuit still dissipates large amounts of power, takes significant area, and has limited scalability.

SUMMARY OF THE INVENTION

In accordance with the present invention, a crosspoint switch comprises a

10 plurality of input buses on which signals are driven at low swing and a plurality of output buses on which signals are driven at low swing. Each of a plurality of crosspoints selectively passes a signal from a low swing input bus to a low swing output bus. Specifically, at the crosspoints, low swing signals on the input buses are sensed, and the signals are driven on the output buses at low swing. Low swing signals on the

15 output buses are then sensed at an output.

In certain embodiments, each crosspoint comprises an amplifier and a low swing driver circuit. The preferred amplifier is a clocked regenerative amplifier. A timing circuit may control timing of the crosspoint switch from a clock. The timing circuit includes a delay, the timing of which varies in a manner similar to timing variations in

20 the driver circuit.

The signals on the input buses and output buses may be differential signals. Low swing drivers at the inputs and at the crosspoints may be push-pull driver circuits, each of which drives a pair of differential lines, one line driven high while the other line is pulled low.

25 BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular description of preferred embodiments of

the invention, as illustrated in the accompanying drawings in which like reference characters refer to the same parts throughout the different views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

5 Figure 1 illustrates a prior high-speed on-chip link circuit using low swing differential signaling.

Figure 2 illustrates the signals of Figure 1.

Figure 3 is an electrical schematic of an improved high speed data link in accordance with the present invention.

10 Figure 4 illustrates a clocked receiver for use in the circuit of Figure 3.

Figure 5 illustrates an alternative clocked receiver for use in the circuit of Figure 3.

Figure 6 illustrates waveforms of signals in the circuit of Figure 3.

Figure 7 illustrates an on-chip data link including a repeater circuit.

15 Figure 8 is a block diagram of the repeater circuit of Figure 7.

Figure 9 is a detailed electrical schematic of the repeater circuit of Figure 8.

Figure 10 illustrates a self-timing circuit for generating the input signals to the driver of Figure 3.

Figure 11 illustrates waveforms of signals in Figure 10.

20 Figure 12 illustrates a self-timing circuit for double data rate operation.

Figure 13 illustrates waveforms of signals in the circuit of Figure 12.

Figure 14 illustrates a self-timing circuit for use in the repeater of Figure 7.

Figure 15 illustrates a double-data-rate self timing circuit for use in the repeater of Figure 7.

25 Figure 16 illustrates the general organization of a crossbar switch.

Figure 17 illustrates details of one row and column bit slice of the crossbar switch of Figure 16 in accordance with the present invention.

Figure 18 illustrates waveforms of signals in the circuit of Figure 17.

Figure 19 illustrates details of the output timing circuit of Figure 17.

Figure 20 is an electrical schematic of the decoder and repeater of Figure 17.

Figure 21 illustrates the predecoder of Figure 17.

Figure 22 illustrates an alternative predecoder for use in the circuit of Figure 17.

Figure 23 illustrates an alternative decoder circuit for use in the circuit of Figure 5 17.

Figure 24 illustrates an alternative embodiment of the crosspoint switch using a pass transistor crosspoint.

Figure 25 illustrates waveforms of signals in the circuit of Figure 24.

DETAILED DESCRIPTION OF THE INVENTION

10 A description of preferred embodiments of the invention follows.

One aspect of the present invention is an on-chip link circuit that moves a digital signal from one point on a chip to another, usually several mm away. The link circuit of the present invention uses a push-pull driver, preferably an H-bridge of NFETs and line-to-line precharge to achieve substantial speed and power advantages compared to 15 prior-art link circuits.

A prior high-speed on-chip link circuit (of the type described in Dally and Poulton, Digital Systems Engineering, pp. 385-387) is illustrated in Figure 1 and waveforms for this prior-art circuit are shown in Figure 2. This circuit reduces the delay and power dissipation of on-chip signaling paths by employing differential signaling 20 with a low-voltage swing using a precharged signaling arrangement. However, while such systems reduce power by lowering the voltage swing, they also consume additional power to precharge the lines each cycle. Also, this precharge operation limits the maximum rate at which the lines can operate.

The driver circuit for this link, 100, consists of a pair of PFET precharge 25 devices, 101 and 102, and a pair of NFET open-drain drivers, 103 and 104. The precharge devices and drivers are coupled to a differential on-chip line 107 that may be many mm in length and thus have significant capacitance. An amplifier 109 senses the

difference in voltage between the two conductors of the differential line 107. Each of the transistors 101-104 forms a leg of an H-bridge.

Timing for this prior circuit is shown in Figure 2. As is typical of precharged circuits, operation of the circuit alternates between precharge intervals, during which 5 both conductors are driven to a common state, and data intervals, during which one conductor is driven to a different state. First, the link conductors xP and xN, 107, are precharged by PFETs 101 and 102 when signal pcN is asserted low. Then to send a binary "1" over the link, pcN is deasserted and signal dP is asserted. This causes conductor xN to discharge by a small amount (typically by 100-300mV). Amplifier 109 10 senses that $V(xP) > V(xN)$ (where $V(xP)$ is the voltage on conductor xP) and thus signals a "1". After sending a "1" the link is again precharged by deasserting dP and asserting pcN. Once precharge is complete, another bit may be sent over the link. In Figure 2, a binary "0" is sent after the second precharge by deasserting pcN to stop the precharge and asserting dN to discharge conductor xP by a small amount. The sense 15 voltage, the amount by which xN and xP must differ to reliably send a bit is set by the noise environment and the offset voltage of receiver 109. Values between 100mV and 300mV are typical.

The prior circuit of Figure 1 is significantly more efficient than using full-swing CMOS logic. However, it leaves considerable room for improvement because of the 20 way it precharges one of the two conductors to a supply rail before discharging one of the conductors through the full sense voltage. For a constant current limit, the present invention essentially doubles the speed and halves the power of the prior art circuit by driving both conductors in opposite directions during the data intervals and precharging to mid-swing by shorting the two conductors together during the precharge intervals. 25 Speed is doubled because each conductor needs to swing only half as far during each interval (half of the sense voltage). With constant current I and constant capacitance C, halving the voltage swing ΔV cuts the time t required in half since charge $Q = C\Delta V = It$. Power is halved as well since the precharge is accomplished by charge sharing between the two lines and thus takes no charge from the power supply. In the prior art circuit,

one of the two conductors is charged from the power supply over the full sense voltage each cycle. With the present invention, one of the two conductors is charged from the power supply through half of the sense voltage each cycle. This takes exactly half the charge, and hence half the power, as the prior art circuit.

5 The improved circuit of the present invention is illustrated in Figure 3. In this circuit, the differential line 107 is driven by an NFET H-bridge driver with shorting precharge 110. Operation of this circuit is illustrated in the waveforms of Figure 6. During precharge intervals, signal pc is asserted high turning on shorting NFET 117. This FET shorts the two conductors of line 107 together, sharing charge between them.

10 Both conductors wind up at the midpoint of the voltage swing. The high conductor drops by passing charge to the low conductor which in turn rises. Thus, at the end of the precharge interval, both conductors of line 107 are at a common voltage. Also, no charge is transferred from the power supply to achieve this precharge.

15 During each data interval, rather than driving just one line as in the prior art circuit, both lines of conductor 107 are driven, one high and one low, in a push-pull action. To send a "1" over the link, as illustrated in the first data interval of Figure 6, dP is asserted which turns on NFETs 111 and 112. Conductor xP is driven high by NFET 111 and conductor xN is driven low by NFET 112. Because both lines are driven, each line need only be driven through half of the sense voltage. For example, if the sense

20 voltage were 200mV, each line would only be driven through 100mV. The charge transferred from the power supply to the high-going line each data cycle is $C\Delta V/2$ (where ΔV is the sense voltage) as compared to $C\Delta V$ in the prior art circuit.

25 After the full sense voltage is developed across the conductors of line 107, amplifier 119 is clocked to sense and amplify this voltage. A clocked amplifier requires less power than a continuous time amplifier and rejects noise by looking only at the wide portion of the eye of the differential signal. The line is then precharged by deasserting dP and reasserting pc to short the two conductors together again.

During the second data interval shown in Figure 6, a "0" is sent across the line by deasserting pc to terminate the precharge and asserting dN which turns on FETs 113

and 114. Conductor xP is pulled low by FET 113 and conductor xN is driven high by FET 114. Again, because both conductors are driven, each conductor need only swing through half of the sense voltage to develop the full sense voltage across the pair of conductors.

- 5 Because of the charge-sharing precharge, the circuit settles to a common-mode voltage that gives equal swings on the high-going line and the low-going line. The common-mode voltage is typically near the negative power supply. For example, for one process, using equal sized pull-up and pull-down devices in driver 110, a voltage supply of 1.8 volts and a 200mV sense voltage, the common mode voltage was 250mV.
- 10 The sense voltage is the difference in voltage between the line in the high state and the line in the low state, and is equal to the swing of a line from low state to high state. Thus, in this example, the low swing of each line is .2 volts compared to a full swing of 1.8 volts or one-ninth of full swing. Low swing is certainly less than one half of a full swing and is usually less than one quarter of full swing.
- 15 Two embodiments of the clocked receiver 119 of Figure 3 are illustrated in Figures 4 and 5. Figure 4 illustrates a receiver that couples the line 107 to the sense nodes of the amplifier via a pair of pass transistors 121 and 122. Figure 5 shows a gate-isolated sense amplifier. Both of these amplifiers have the property that one output is asserted (high) when the clock is asserted (low) and both outputs are deasserted (low)
- 20 when the clock is deasserted (high).

For the pass-transistor-isolated amplifier, the outputs follow the inputs through pass transistors 121, 122 when the clock is high. Typically, a latch circuit follows the amplifier and the swing of the oP and oN signals is insufficient to activate the latch. Being below the required logic levels, the signal looks like a logic 0 to the last circuit.

- 25 However, when the circuit is clocked by a low clock signal ckN, the pass transistors 121 and 122 turn off to capture the signal on lines oP and oN, and transistor 125 turns on to energize the latch circuit of transistors 127, 128, 131 and 132. Assume that the signal on line oP is higher than the signal on line oN. The signal on line oP tends to turn transistor 131 off and to turn transistor 128 on, and the signal on line oN tends to turn

the transistor 132 on and transistor 127 off. The current through transistor 132 drives oP higher to fully turn on transistor 128 and turn off transistor 131. With transistor 128 fully on, signal oN is pulled to ground, fully turning off transistor 127 and turning on transistor 132. As a result, the partial swing received on lines xP and xN is sensed and 5 driven through a full swing on lines oP and oN when the clock is asserted (low).

For the gate isolated amplifier of Figure 5, the outputs are predischarged to GND when the clock is high through transistors 149 and 150. When ckN is asserted low, transistors 149 and 150 are turned off and voltage is applied through transistor 135 to the sense amplifier. Assume that xP is higher than xN in the received partial swing 10 differential signal. Transistor 137 tends to turn off as transistor 138 tends to turn on. With both gates and sources of transistors 141 and 142 at ground, transistor 142 begins to conduct. With charge building on oP, transistor 141 is driven fully off and transistor 145 is driven fully on to pull signal oN full low. Signal oN similarly drives transistor 142 fully on and transistor 146 fully off to drive output oP to the supply voltage level. 15 Thus the partial swing on signals xP and xN is sensed to a full swing differential signal on lines oP and oN.

Both the pass-gate input amplifier of Figure 4 and the gate isolated amplifier of Figure 5 are examples of *regenerative* amplifiers in that, when enabled by a clock, they derive their gain by using positive feedback to reinforce a small signal. In both 20 amplifiers there is originally a small signal 0.2V or less across output nodes oP and oN. This signal is amplified by the positive feedback of a pair of back-to-back inverters. In Figure 4 the inverters are formed by transistors 127, 128, 131, and 132. In Figure 5 transistors 141, 145, 142, and 146 form the inverters. In either case, once the inverters are enabled, they increase the magnitude of the voltage across the output nodes oP and 25 oN exponentially, increasing the voltage by a factor of e (2.718) each time constant until oP and oN reach the power rails. A typical regenerative amplifier in an 0.18 μ m CMOS process has a time constant of about 100ps, so a regeneration time of just 1 ns gives a gain of e^{10} , over 22,000.

The circuit of Figure 5 is the sense amplifier which is now used. Both of the amplifier configurations of Figures 4 and 5 are well known and many other sense amplifier configurations could be used.

An on-chip wire is fundamentally an RC transmission line and hence has a delay
5 that increases quadratically with length, since both R and C increase linearly with length. For this reasons, it is advantageous to break long on-chip lines into multiple segments separated by repeaters. In the prior art such repeaters are most often realized as CMOS inverters that require a full-swing input and produce a full-swing output. Such full-swing CMOS repeaters suffer the same power and delay disadvantages as
10 full-swing drivers and receivers.

The present invention can be used to realize an on-chip transmission line with one or more repeaters as illustrated in Figure 7 for the case of a single repeater. The line 107 is broken into two segments, 107A and 107B. During each data interval, as signaled by ck1 being asserted, driver 110 drives the current bit onto line segment 107A
15 in exactly the manner described in relation to Figures 3 and 6. After line 107A reaches the sense voltage, ck2 is asserted. This triggers repeater 160 to sense the voltage on segment 107A, determine whether the current bit is a "1" or a "0", and repeat this bit by driving segment 107B. After segment 107B is driven to the sense voltage, ck3 is asserted which triggers amplifier 119 to sense the voltage on this segment and generate
20 the output for the entire line.

Because this transmission line is on chip, the clock supplied to the driver 110, received 119 and repeater 160 can be precisely timed relative to each other from a common clock for both low power and high speed operation. Each circuit is clocked by one of plural clock signals, having different phases, which are provided by the common
25 clock. Prior off-chip repeater circuits which recover the clock from the signal and regenerate the signal using that clock have lacked the advantage of the use of a precise on-chip clock.

Repeater 160 combines the functionality of receive amplifier 119 and driver 110. In fact, as illustrated in Figures 8 and 9, repeater 160 can be realized by combining a

receive amplifier 119 and a driver 110 back to back. Because the outputs of receive amplifier 119 are deasserted when the clock is inactive (high), these signals can be used directly to drive the dP and dN inputs of driver 110. When the clock is asserted, the receive amplifier drives the correct driver input. When the clock is deasserted (high),

5 the both driver inputs dP and dN are deasserted and the driver is precharged by connecting its precharge input to the clock.

The width of the dP and dN pulses into driver 110 determines the voltage swing on line 107 (or segments 107A and 107B). To ensure that the voltage swing is large enough to be reliably sensed but not so large as to dissipate excess power, the preferred

10 embodiment uses a self-timing circuit to control the width of the pulses.

A self-timing circuit for controlling the width of the driver input pulses is illustrated in Figure 10 and waveforms for this circuit are shown in Figure 11. When input clock ck1 is asserted, the clock enables AND gates 163 and 164 to assert either dP or dN depending on the state of input d. After a delay of t_w , signal ckd, the output of

15 delay line 165 is asserted disabling both dP and dN. Thus, the pulse generated on dP or dN will have a width of t_w . The ckd signal also causes the precharge, pc, signal to be asserted via OR gate 167. This causes driver 110 to be precharged whenever a data input is not asserted. An inverted version of the precharge signal, ck2N, generated by inverter 166, can be used to clock receiver 119 or repeater 160.

20 Preferably, delay line 165 is designed so that its delay tracks process, voltage, and temperature variations in the same manner as driver 110 so that while variation of these parameters will cause pulse width variation, the voltage swing on line 107 will remain constant across process voltage and temperature. A faster driver, for example, will charge to a higher voltage over a given pulse duration, but a faster delay circuit

25 which tracks the speed of the driver will offset the speed of the driver with a shorter pulse duration. This tracking can be accomplished using methods well known in the art. For example, delay line 165 can be realized from an even number of inverters with each inverter pair loaded by a wire load fabricated on the same layer as line 107. With

appropriate ratios of devices, the delay of this line will closely track the delay of driver 110 driving line 107.

Because the circuit of the present invention is capable of very high speed, in some applications it is advantageous to drive data on both edges of the clock signal.

- 5 Such 'double-data-rate' operation can be accomplished using the self-timing circuit shown in Figure 12. The operation of this circuit is best understood by referring to the waveforms shown in Figure 13. As with the circuit of Figure 10, delay line 165 with delay t_w controls the pulse width of the asserted data signals. Unlike Figure 10, however, the circuit of Figure 11 asserts data signals on both edges of the clock.
- 10 AND-OR gates 171 and 172 multiplex data inputs d1 and d0 to provide data from d1 when ck1 is high and ckd is low, and d0 when ck1 is low and ckd is high. Data is provided by asserting output dP high if the selected data input is true (high) or asserting dN high if the selected data input is false (low). When ck1 and ckd are in the same state, high or low, both dP and dN are low and the precharge signal, pc, is asserted.

- 15 When sending bits on both edges of the clock over line 107, the inverted version of the pc signal, ck2N, can be used to clock the data through a repeater 160. Alternatively, data may be demultiplexed using a pair of receivers 119 by clocking one receiver on the rising edge of ckd (to recover d1) and clocking a second receiver on the falling edge of ckd (to recover d0).

- 20 When a repeater is used as illustrated in Figure 7, the pulse width of the driving signals internal to the repeater can be self-timed as well by using the circuit of Figure 14 (for sending data on one edge of the clock), or Figure 15 (for sending data on both edges of the clock). These circuits delay signal ckd by t_w using delay line 181 to generate delayed clock signal cke. For the single-edged case (Figure 14), NAND gate 182 asserts ck2N low, enabling repeater 160 to drive segment 107B, during the t_w long period when ckd is high and cke is low. For the double-edged case (Figure 15), XNOR gate 185 asserts ck2N low during both t_w long periods when ckd and cke are in opposite states. When the repeater is used, a receive amplifier or a next-stage repeater can either
- 25

be clocked on the rising edge of ck2N or, for the double-edge case, on both edges of cke.

One skilled in the art will understand that the logic gates in Figures 10, 12, 14, and 15 may be realized as static CMOS gate circuits or using other well-known circuit techniques such as domino logic or pass-transistor logic.

High-Performance, Low-Power Switch Core

A second aspect of the present invention adapts the low-power, high-speed signaling circuit described above to the design of a high-speed CMOS crossbar switch core. The advantages of a low swing data link are multiplied by the large quantity of links through the many crosspoints of a large crossbar switch. The large power requirements of full swing data links in a very large crossbar switch would run the risk of collapsing the power supply.

Figure 16 shows the general organization of a crossbar switch. The preferred embodiment of the present invention is a 72 x 72 synchronous CMOS crossbar switch that connects any of the 72 STS-96 (4.976Gb/s) input channels to each STS-96 output channel and can be reconfigured on each byte. The switch operates at a data rate of 622MHz and achieves STS-96 bandwidth by switching one byte (8 bits) of data during each data period. To simplify timing, the switch uses a 311MHz clock and switches data on both edges of the clock. The aggregate input bandwidth of the switch is 358Gb/s. The output bandwidth is the same for a total bandwidth (input plus output) of 716Gb/s.

The crossbar switch includes 72 8-bit wide input buses 191, 72 output address select buses 193, and 72 8-bit wide output buses 195. Each input bus is connected to each output bus by an eight-bit wide crosspoint, xp 197. There are 5,184 byte-wide crosspoints 197 in the complete switch that are composed of 41,472 individual single-bit crosspoints.

The preferred embodiment is realized using differential signaling so each signal line in Figure 16 uses two conductors, one to carry the positive rail and one to carry the

negative rail of the differential signal. Thus the 72 8-bit wide input signals 191 comprise 576 signals carried on 1152 wires.

During each 1.6ns data period, a byte of data is driven onto each of the 8-bit differential input buses 191. At the same time an address, selecting one of the 72 input 5 buses, is encoded on each of the 72 address buses 193. Each address bus selects the input bus that is to be connected to the corresponding output bus by enabling the crosspoint 197 connected to that bus. At most one crosspoint in each column is enabled. In addition to point-to-point communications, the switch can be configured to perform an arbitrary multicast or broadcast operation by having multiple address buses 10 193 select the same input bus 191.

Each crosspoint that is enabled samples the data on the input bus it is connected to and drives this data on the output bus it is connected to in a manner identical to the repeater of Figures 8 and 9. The output data is then sampled off the output bus by a receiver (not shown). After the input buses are sampled, each of the 576 input bits are 15 precharged to mid-swing by shorting the two rails of each bit together. The output buses are precharged in a similar manner after they are sampled.

Figure 17 shows additional details of one row and column bit slice of the crossbar switch core of Figure 16. That is, the figure shows one bit 191A of one of the input buses 191 and one bit 195A of one of the output buses 195 along with all of the 20 corresponding address bus 193. Signals in Figure 17 associated with each row, such as input bit 191A are labeled with the Y-coordinate of the row; for example, the two rails of input bit 191A are labeled hYP and hYN. Similarly, signals associated with each column are labeled with the X-coordinate of the column, such as addrX, vXp, and vXN. Signals associated with a crosspoint are labeled with both the X-coordinate of the 25 column and the Y-coordinate of the row such as ckXYN. Thus, Figure 17 depicts the circuitry associated with one bit of row Y and one bit of column X of the crossbar switch of Figure 16. One skilled in the art will understand how to repeat the circuitry in this figure to populate the entire switch.

The operation of the crossbar bit slice shown in Figure 17 is best understood by referring to the waveforms shown in Figure 18. Each time the input clock, ck, toggles, the double-data rate input timing circuit 177 asserts a drive signal dN or dP to drive the selected data bit, d1 on clock high and d0 on clock low, onto differential horizontal input line hYP, hYN 191A via driver 110. When the sense voltage has been developed across this line, the timing circuit 177 deasserts dN and dP and asserts precharge signal, pc, and toggles delayed clock, ckd. The precharge signal causes driver 110 to initiate precharge of the horizontal input line 191A by shorting the two polarities of this line together. Toggling delayed clock ckd causes the output timing circuit 203 to deassert the vertical precharge, vpc 199. Deasserting vpc serves both to disable precharge device 215 and to clock crosspoint decoder 201. When the vpc signal 199 is deasserted and the row address encoded on address bus addrX 193 matches the Y-coordinate of the current row, the crosspoint clock ckXYN 219 is asserted. Asserting crosspoint clock 219 causes the modified repeater 217 in the crosspoint to sample input line 191A and to drive column X vertical output line vXP,vXN 195A. Note that this sampling and drive only occurs in crosspoints that are selected by the predecoded address on addrX 193. Also, unselected crosspoints do not precharge the vertical line 195A because the repeater is modified from the circuit shown in Figure 9 as shown in Figure 20 to omit precharge transistor 117 from the driver 223. Instead this precharge is performed for the entire column by precharge device 215. Thus the crosspoint clock, ckXYN 219 serves both as a clock and as an output enable.

After the modified repeater 217 in selected crosspoint 197A has driven the vertical line vXP, vXN 195A through the sense voltage, output timing circuit 203 toggles delayed clock cke and asserts the vertical precharge, vpc 199. Toggling clock 25 cke samples the bit on the vertical line 195A into one of the receive amplifiers 207 and 208. When cke goes high, the data is sampled into amplifier 207 to recover on signal q1 the signal clocked in on d1 of the selected input row when ck went high. Similarly, when cke goes low, the data on the vertical line 195A is sampled into amplifier 208 to recover on signal q0 the data clocked in on d0 of the selected input row when ck went

low. Receivers 207 and 208 are identical in design to receiver 119 shown in Figure 5 except that the clock to receiver 208 is inverted.

The crossbar is reconfigured each byte time, selecting a new input row to be connected to each output column under control of the address lines 193. Each cycle, 5 two 7-bit addresses, a1X and a0X, are input to column X of the crossbar. Address a1X selects one of 72 input rows to be connected to output column X when the clock is high while a0X selects the input row to be selected when the clock is low. These two addresses are multiplexed by multiplexer 211 under control of delayed clock ckd. When ckd is low, a1X is selected and when ckd is high, a0X is selected. The output of 10 the multiplexer is pre-decoded from a 7-bit binary format into an 8-bit one-hot plus a 9-bit one hot encoding. The resulting 17 lines are distributed up the column as addrX 193. The falling edge of vpc samples the relevant bits of addrX into the final decoder 201 associated with each crosspoint. This sampling occurs before addrX can change in response to ckd switching the multiplexer. Final decoder 201 samples one bit from the 15 8-bit address group and one bit from the 9-bit address group. If both sampled bits are high, the current crosspoint is selected and crosspoint clock ckXYN 219 is asserted to enable the modified repeater 217 to drive the vertical lines 195A.

The input timing circuit 177 has already been described with respect to Figure 12. Details of output timing block 203 are shown in Figure 19. Delayed clock ckd is 20 further delayed by delay line 181 to generate clock cke. Delay line 181 has a delay tw that is matched to the time required for the modified repeater 217 to drive the vertical lines 195A through the sense voltage. This matching is achieved in the same manner as described above in reference to Figures 10 through 15. Clocks ckd and cke are combined by XNOR gate 221 to generate the vertical precharge signal vpc 199. This 25 signal is asserted from the time cke toggles to the time ckd toggles in the opposite direction.

Figure 20 shows the details of the entire crosspoint bitslice 197A including final decoder 201 and modified repeater 217. Decoder 201 comprises two CMOS transmission gates 235 and 236, an inverter 233, a buffer 237 and a 2-input dynamic

NAND gate comprising transistors 225-228. When vertical precharge clock vpc 199 falls, it turns off transmission gates 235 and 236 latching address bits aXai and aXbj onto the gates of transistors 225 and 226 respectively. An inverter delay later, the output of inverter 233 goes high triggering the dynamic NAND gate 225-228. The 5 output of the NAND gate goes low only if both aXai and aXbj were high at the point they were sampled by vpc falling. Each row Y is connected to one of the eight aXa lines and one of the nine aXb lines so that $Y = 8j + i$. Thus, the coincidence of aXai and aXbj indicates that the current row, Y is selected. The output of the NAND gate is buffered by buffer 237, a pair of inverters, to generate the crosspoint clock ckXYN 219. When 10 the current row, Y, is selected, crosspoint clock 219 falls a short time after the falling edge of vpc 199. If the row is not selected, the crosspoint clock remains high.

Modified repeater 217 comprises receiver 119 and modified driver 213. This repeater is identical to the repeater 160 of Figures 8 and 9 except that the H-bridge driver 223 has no precharge device 117. When a crosspoint is selected and the 15 crosspoint clock ckXYN 219 falls, the receiver amplifies the signal on differential input hYP/hYN and asserts one of its two outputs oP or oN high depending on the polarity of the input. The asserted output enables the H-bridge driver 223 to drive the vertical lines vXP/vXN in the proper direction. When vpc goes high, clock ckXYN goes high as well, predischarging receiver 119 via transistors 149 and 150. Discharging causes 20 receiver outputs oP and oN to both go low disabling driver 223. When a crosspoint is not selected, ckXYN 219 remains high causing receiver outputs oP and oN to remain low which in turn disables driver 223 so some other crosspoint can drive the vertical lines without interference.

Figure 21 shows the details of predecoder 205. The seven-bit address ax 245 25 from multiplexer 211 which carries binary numbers in the range from 0 to 72 is split into a 3-bit field 247, the least significant bits, that is encoded in binary with a range from 0 to 7 and a four-bit field 248, the most significant bits, also binary with a range from 0 to 8. Three-bit binary address field 247 is fed to a 3 to 8 decoder 241 which decodes the eight input possibilities into an eight-bit one-hot encoding on signal

aXa[7:0] 251. If field 247 contains the binary encoding of i, signal aXai will be true and all other bits of aXa will be false. Similarly, four-bit field 248 is decoded from binary to one-of-nine encoding by 4 to 9 decoder 243. The nine-bit one-hot encoding is output on signal aXb[8:0] 252. If field 248 contains the binary encoding of j, signal 5 aXbj will be true and all other bits of aXb will be false. Collectively the nine bits of aXb[8:0] and eight bits of aXa[7:0] make up the seventeen bit pre-decoded address bus addrX 193. One skilled in the art will understand how decoders 241 and 243 can be constructed using static CMOS gates, dynamic CMOS gates, or pass-transistor logic.

One skilled in the art of decoder design will also understand that the seven input 10 bits could be divided differently for predecoding, for example into 2, 2, and 3 bit fields rather than the 3 and 4 bit fields shown here. While the 2, 2, and 3 bit division results in fewer lines of predecoded address in bus addrX 193, only 13 rather than 17, it dissipates more power since three predecoded address lines may toggle each cycle rather than two.

15 An alternative form of predecoder 205, labeled 205A, that combines the vertical precharge clock 199 with the predecoded address bits aXa[7:0] 251 is illustrated in Figure 22. This predecoder is used in conjunction with the modified final decoder 201A shown in Figure 23. The predecoder 205A is identical to predecoder 205 shown in Figure 21 except for two changes. First, the low-three bits 247 of address aX 245 are 20 latched by three-bit transparent latch 255 to hold them valid after vpc 199 goes low and ckd switches multiplexer 211. Latch 255 essentially replaces the latch formed by CMOS transmission gate 235 in Figure 20 by moving this latching function from after 3 to 8 decoder 257 to before the decoder. Second, a 3 to 8 decoder with a low-true enable 257 replaces the simple decoder 241 of Figure 21. With this decoder, all eight bits of 25 aXa[7:0] 251 are low whenever vpc 199 is high. When vpc 199 goes low, one bit, aXai, corresponding to the binary value i on decoder input 249, of 251 goes high while the remaining bits stay low. Thus signals aXa[7:0] 251 are qualified clocks that combine the timing of vpc 199 with the predecoding of the low three bits of the address 249.

The alternative predecoder 205A of Figure 22 is used in conjunction with the modified final decoder 201A shown in Figure 23. When the current crosspoint, XY, is selected, signal aXbj will be high and signal aXai will pulse high where the current row address $Y = 8j + i$. The qualified clock aXai rises when vpc 199 falls and falls when 5 vpc 199 rises. The rising edge of aXai latches the value of aXbj on the gate of transistor 242 via transmission gate 245. This rising edge also serves to trigger the dynamic NAND gate formed by transistors 241-243 to pull the crosspoint clock ckXYN 219 low via buffer 235. If the current crosspoint is not selected, either there will be no pulse on aXai, in which case the NAND gate is never triggered, or aXbj will be low, in which 10 case the NAND gate evaluates to a high value and ckXNY 219 is not asserted. When aXai returns low in the all crosspoints, all of the dynamic NAND gates are precharged via device 243.

An alternative embodiment of the present invention that uses a passive pass-transistor crosspoint rather than a regenerative repeater crosspoint is illustrated in 15 Figure 24. Waveforms showing the operation of this embodiment are shown in Figure 25. Each time ck toggles, the input timing circuit 177A deasserts precharge clock pc and asserts one of the drive signals dP or dN to drive the current bit onto the horizontal lines hYP/hYN 191A. At the same time, pc going low enables decoder 201B to decode the address encoded on lines addrX 193. If the address matches the current row, Y, then 20 signal enXY 263 is asserted high. Decoder 201B differs from decoder 201 in that non-inverting buffer 235 in Figure 20 or 23 is replaced by an inverting buffer to give a high true output. When enable signal enXY 263 is asserted, it enables NMOS pass-transistors 261 and 262 to pass the signal from the horizontal lines 191A to the vertical lines 195A.

25 When the vertical lines have been driven through the sense voltage, the input timing circuit 177A deasserts the drive signals dP and dN, asserts the precharge signal, pc, and toggles delayed clock ckd. Clock ckd samples the value on the vertical lines into one of the receive amplifiers 207 or 208 and selects the next address via

multiplexer 211. Asserting the pc signal precharges both the horizontal lines, via driver 110, and the vertical lines, via shorting transistor 215.

The input timing circuit 177A differs from input timing circuit 177 in Figures 12 and 17 only in that delay line 165 has its delay t_w set to match the combined delay 5 across the horizontal lines 191A, pass gates 261 and 262, and vertical lines 195A instead of just the delay of the horizontal lines as in Figure 17. Thus clock ckd is asserted when the sense voltage is developed across the vertical lines.

A disadvantage of the pass-transistor crosspoint of Figure 24, relative to the 10 amplifier crosspoints of Figure 17, is that the pass-transistor crosspoint provides no gain and hence is unable to support fan-out of the low swing signals from a single horizontal line to multiple vertical lines. Such fanout is required for multicast or broadcast 15 operation.

Yet another alternative embodiment of the present invention employs static 15 (unclocked) amplifiers at each crosspoint and at the outputs to enable the crosspoint switch to be operated without a clock. To save power, only the amplifiers at selected crosspoints are powered up. Amplifiers at unselected crosspoints are unpowered which leaves their outputs in a high-impedance state. In the preferred form of this 20 embodiment, each crosspoint consists of an amplifier, a pulse generator, and an H-bridge driver. When the crosspoint is selected, the amplifier amplifies the low swing signal on the horizontal lines to full swing. The pulse generator detects transitions on the amplified signal and generates pulses to drive the H-bridge driver using a self-timed circuit of the form shown in Figure 12. The H-bridge driver then drives the vertical line to the proper level.

One skilled in the art of digital circuit design will understand that several 25 variations of the design described above are possible. For example, the crossbar circuits of Figures 17 and 24 can be modified to operate on only a single edge of the clock rather than on both edges by using the circuit techniques described in conjunction with Figures 10 and 14. The crossbar circuits could also be realized with the prior driver arrangement shown in Figure 1 rather than the driver of Figure 3. The crossbar switch

can be realized with a different number of input and output ports and the number of inputs and outputs need not be the same. Also, the width of each port could be varied and the clock rate of the crossbar could be varied. Further the entire crossbar need not be run in a single timing domain, but rather each input can run in its own clock domain

While this invention has been particularly shown and described with references to preferred embodiments thereof, it will be understood by those skilled in the art that various changes in form and details may be made therein without departing from the scope of the invention encompassed by the appended claims.

卷之三

CLAIMS

What is claimed is:

1. A crosspoint switch comprising:
 - a plurality of input buses, signals on the input buses being driven at low swing;
 - a plurality of output buses, signals on the output buses being driven at low swing; and
 - a plurality of crosspoints, each selectively passing a signal from a low swing input bus to a low swing output bus.
- 10 2. A crosspoint switch as claimed in claim 1 wherein each crosspoint comprises an amplifier.
3. A crosspoint switch as claimed in claim 2 wherein each crosspoint comprises a low swing driver circuit.
4. A crosspoint switch as claimed in claim 3 wherein the amplifier is a clocked regenerative amplifier.
- 15 5. A crosspoint switch as claimed in claim 4 further comprising a timing circuit which controls timing of the crosspoint switch from a clock, the timing circuit including a delay, the timing of which varies in a manner similar to timing variations in the driver circuit.
- 20 6. A crosspoint switch as claimed in claim 3 wherein the signals on the input buses and the output buses are differential signals.

7. A crosspoint switch as claimed in claim 6 wherein low swing drivers which drive the input buses and the low swing drivers at the crosspoints are push-pull driver circuits, each of which drives a pair of differential lines, one line driven high while the other line is pulled low.

5 8. A crosspoint switch as claimed in claim 1 further comprising a plurality of amplifiers which amplify the signals on the output buses, the amplifiers being clocked regenerative amplifiers.

9. A crosspoint switch as claimed in claim 8 wherein the signals on the input buses and the output buses are differential signals.

10 10. A crosspoint switch comprising:
a plurality of input buses;
a plurality of low swing drivers which drive signals to the input buses, each low swing driver driving a pair of differential lines, one line driven high while the other line is pulled low;

15 a plurality of output buses;
a plurality of crosspoints, each selectively passing a signal from an input bus to an output bus, each crosspoint comprising an amplifier which amplifies a signal on an input bus and a low swing driver which drives a low swing signal on an output bus; and

20 a plurality of output amplifiers which sense the signals on the output buses.

11. A crosspoint switch as claimed in claim 10 further comprising a timing circuit which controls timing of the crosspoint switch from a clock, the timing circuit including a delay, the timing of which varies in a manner similar to timing variations in the driver circuit.

25

12. A crosspoint switch as claimed in claim 11 wherein the amplifier is a clocked regenerative amplifier.
13. A method of connecting signals from a plurality of input buses to a plurality of output buses comprising:
 - 5 driving signals on the input buses with a low swing;
 - at crosspoints between the input buses and output buses sensing signals on the input buses and driving signals on the output buses at low swing; and
 - sensing the low swing signals on the output buses.
14. A method as claimed in claim 13 wherein the signals are sensed by a clocked regenerative amplifier.
- 10
15. A method as claimed in claim 14 further comprising controlling timing of the crosspoint switch from a clock through a timing circuit including a delay, the timing of which varies in a manner similar to timing variations in driver circuits which drive the signals.
- 15
16. A method as claimed in claim 13 wherein the signals on the input buses and the output buses are differential signals.
17. A method as claimed in claim 16 wherein the signals on the input buses and the output buses are driven by push-pull driver circuits, each of which drives a pair of differential lines, one line driven high while the other line is pulled low.
- 20
18. A method as claimed in claim 13 further comprising amplifying the signals on the output buses in amplifiers, the amplifiers being clocked regenerative amplifiers.

19. A method as claimed in claim 18 wherein the signals on the input buses and the output buses are differential signals.
20. A method of connecting signals from a plurality of input buses to a plurality of output buses comprising:
 - 5 driving signals on the input buses through a plurality of low swing drivers, each low swing driver driving a pair of differential lines, one line driven high while the other line is pulled low;
 - 10 at a plurality of crosspoints, sensing signals from the input buses with amplifiers which amplify signals on the input buses, and driving signals on the output buses with low swing drivers; and
 - 15 sensing the low swing signals on the output buses with output amplifiers.
21. A crosspoint switch comprising:
 - means for driving a plurality of low swing signals on a plurality of input buses;
 - 15 a plurality of crosspoint means for sensing signals on the input buses and driving low swing signals on a plurality of output buses.

HIGH-SPEED, LOW-POWER CROSSBAR SWITCH

ABSTRACT OF THE DISCLOSURE

In a crosspoint switch, both input buses and output buses are driven at low swing. Self-timed, differential, push-pull, low swing driver circuits drive the input buses and are provided in the crosspoints to drive the output buses. Clocked, regenerative sense amplifiers are provided in crosspoints and at outputs of the data buses.

卷之三

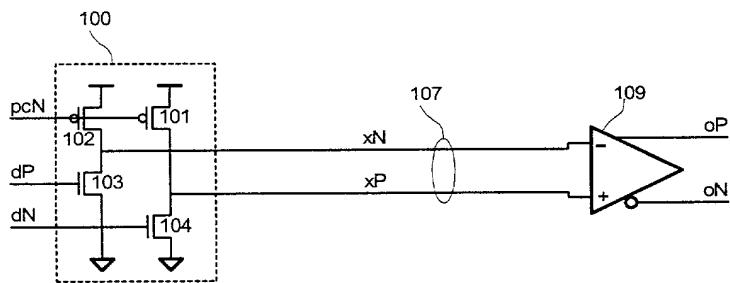


Figure 1: Prior Art

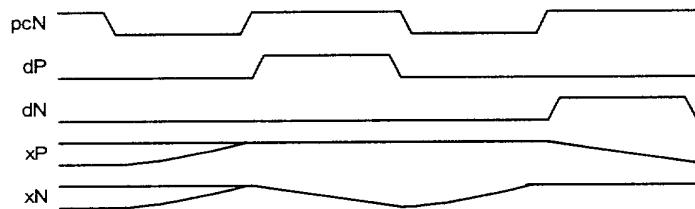


Figure 2: Prior Art

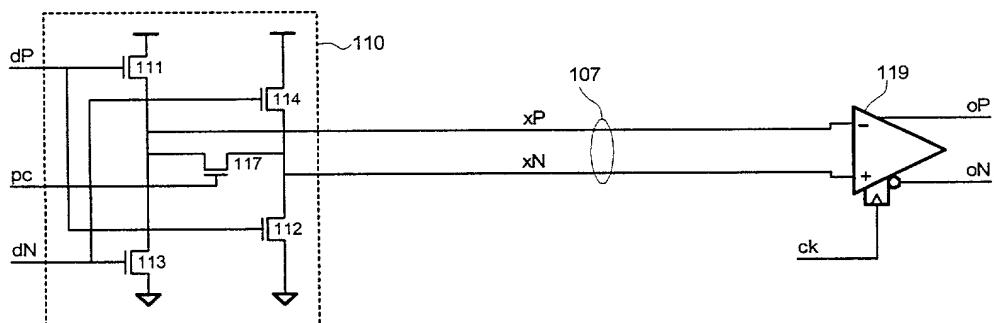


Figure 3

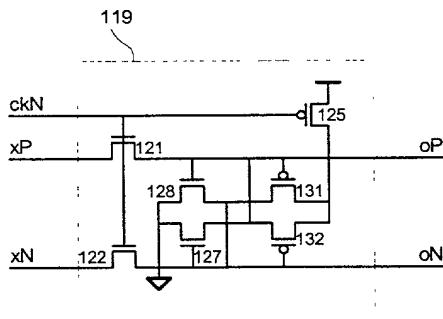


Figure 4

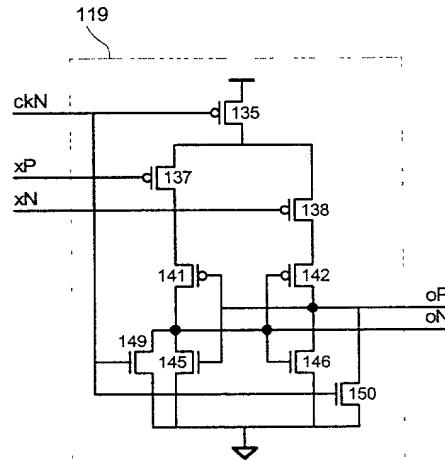


Figure 5

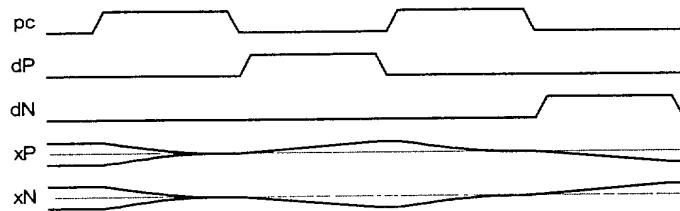


Figure 6

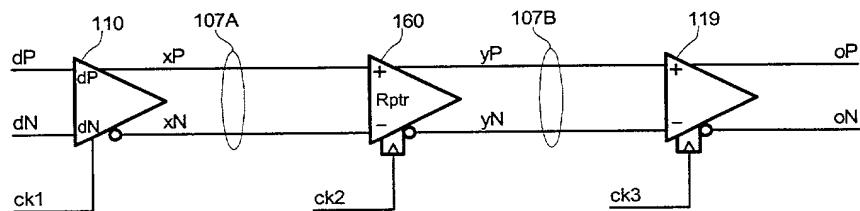


Figure 7

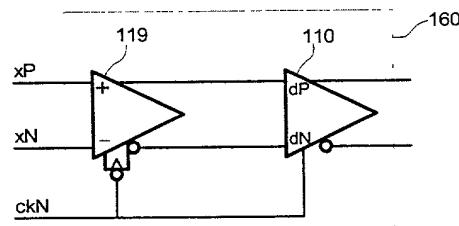


Figure 8

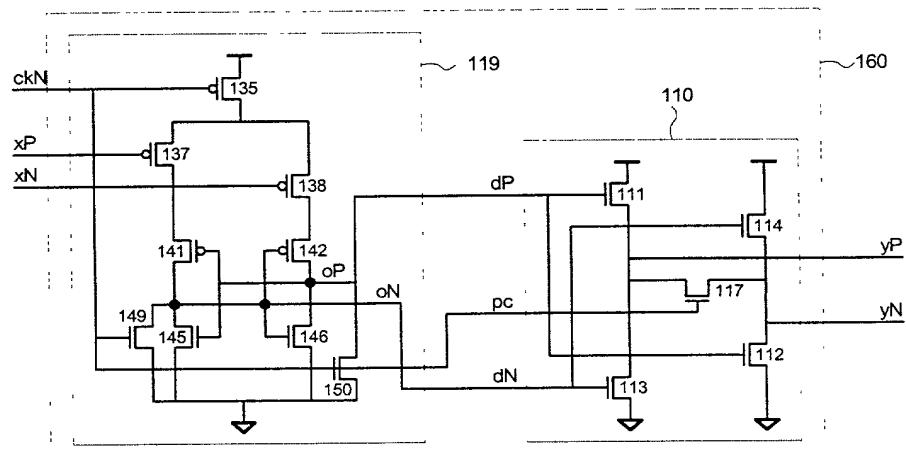


Figure 9

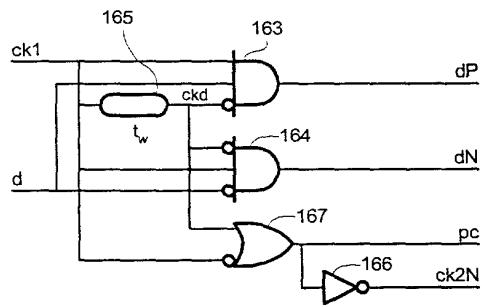


Figure 10

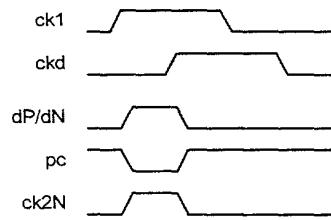


Figure 11

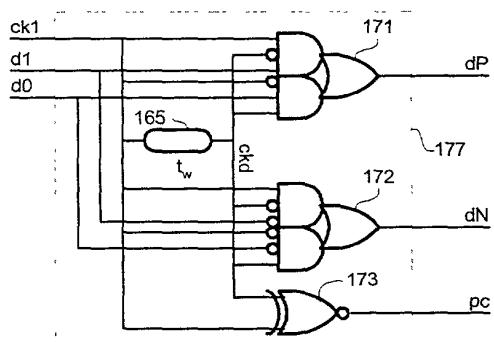


Figure 12

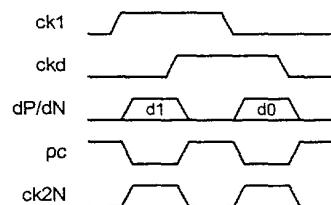


Figure 13

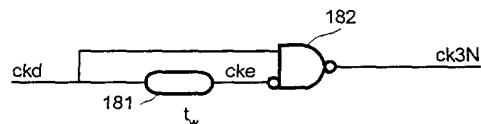


Figure 14

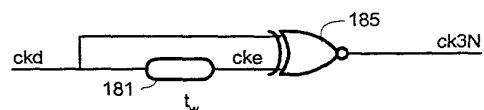


Figure 15

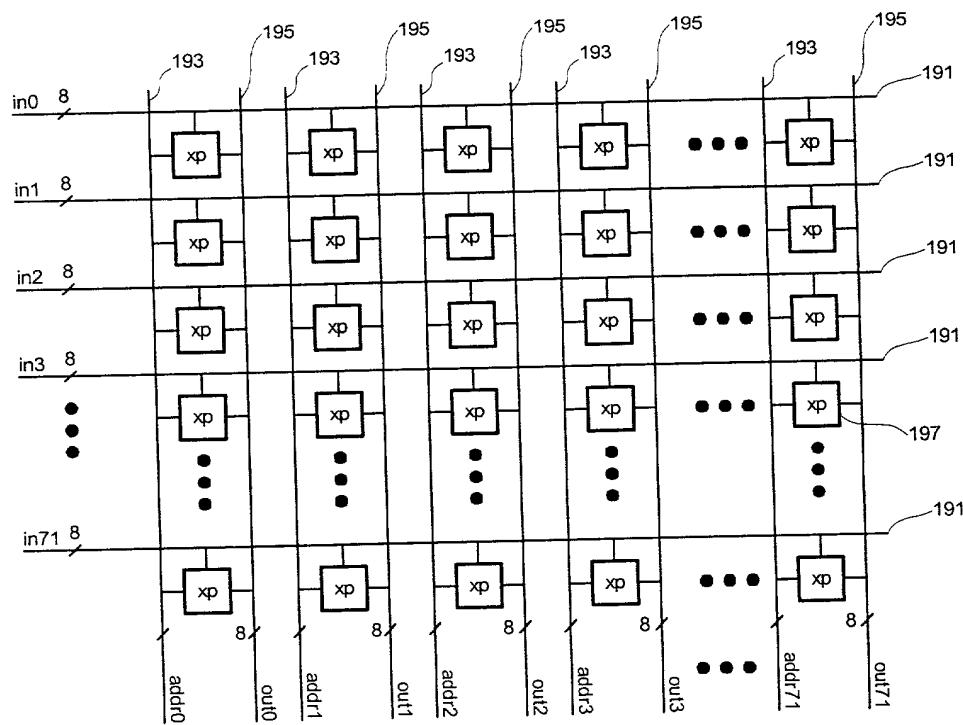


Figure 16

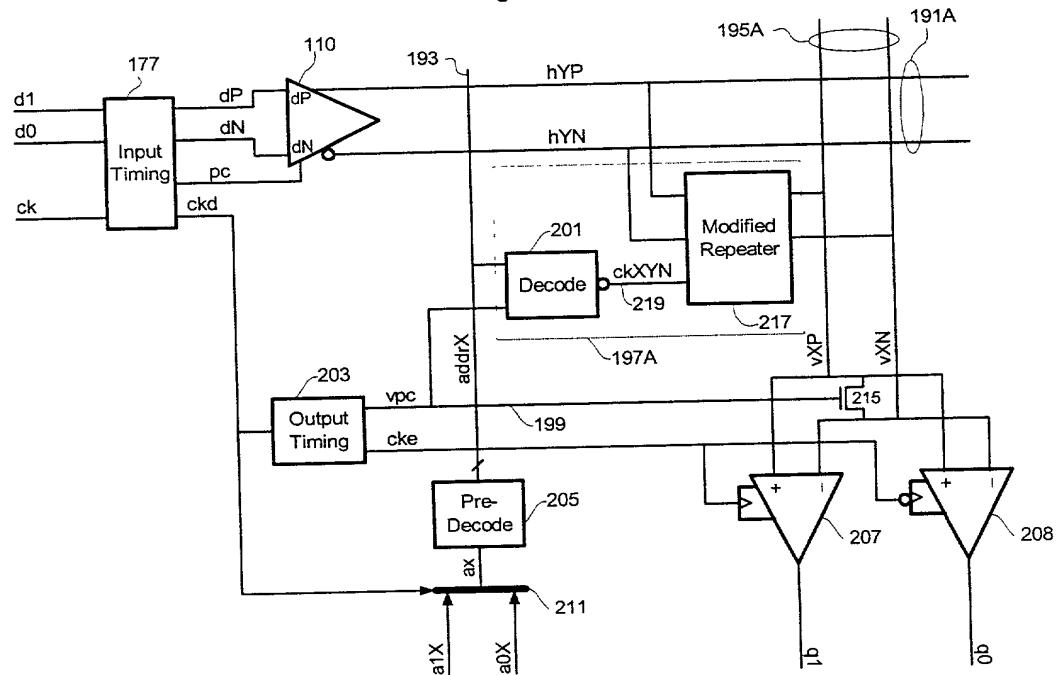


Figure 17

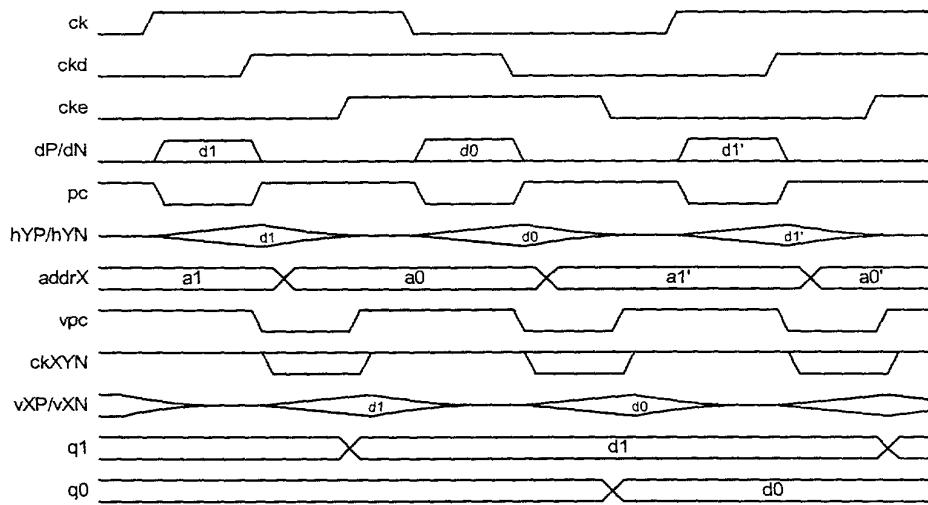


Figure 18

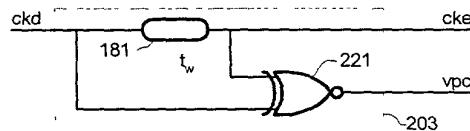
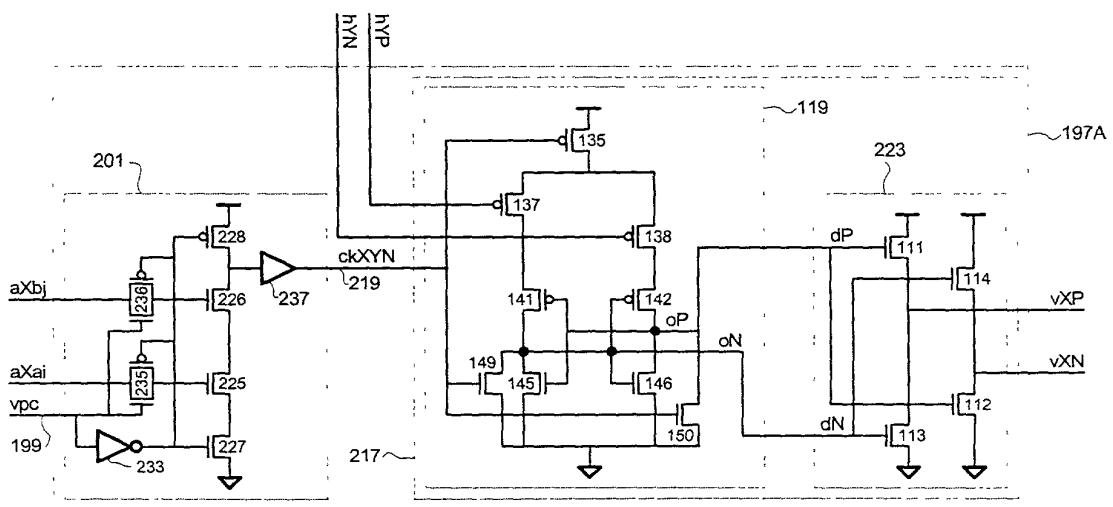


Figure 19



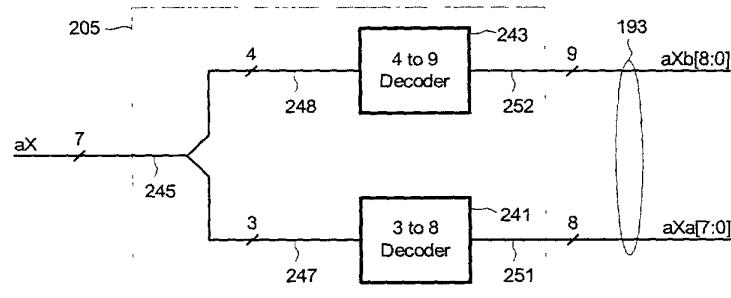


Figure 21

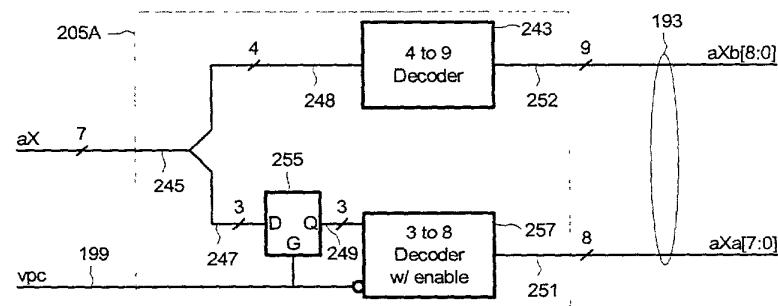


Figure 22

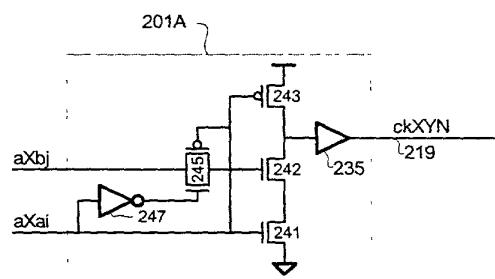


Figure 23

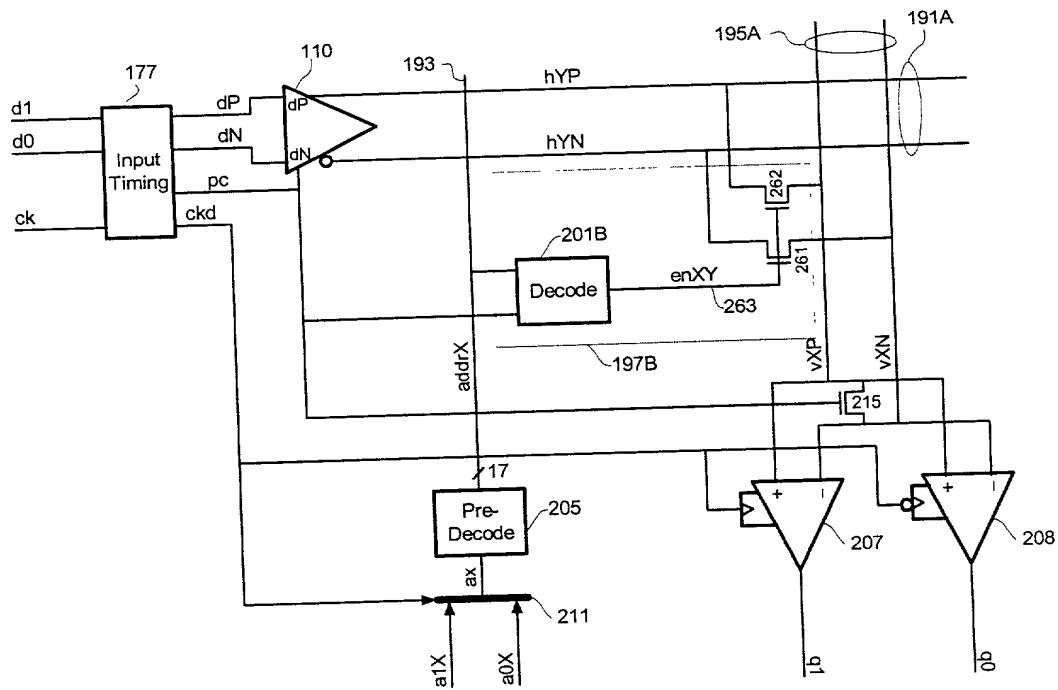


Figure 24

Figure 24 is a block diagram of a digital logic circuit, likely a memory controller or interface. The diagram shows an 'Input Timing' block (177) receiving data d1, d0, and control signals ck, ckd, pc. It outputs dP, dN, and a clocked signal to a buffer (110). The buffer (110) outputs 193 to a 'hYP' block. The 'hYP' block also receives address 'addrX' and control signals 'enXY' (from a 'Decode' block 201B) and 'vXP/vXN' (from a 'Pre-Decode' block 205). The 'hYP' block outputs 195A to a '191A' block. The '191A' block is connected to a 'vXN' block. The 'vXN' block has two outputs: one through a summing junction (207) to 'q1', and another through a summing junction (208) to 'q0'. The 'vXN' block also receives 'vXP' and 'vXN' signals. A 'Pre-Decode' block (205) receives 'addrX' and 'ax' (from the 'hYP' block) and outputs 'a1X' and 'a0X' to a '211' block. The '211' block outputs to 'd1' and 'd0' lines. The 'hYN' block receives 'addrX' and 'enXY' and outputs 'd1'' and 'd0''. The 'hYN' block also receives 'vXP/vXN' and 'vXN' signals. The 'hYN' block outputs '197B' to the 'hYP' block. The 'hYP' block also receives 'vXP/vXN' and 'vXN' signals. The 'hYP' block outputs '193' to the buffer (110).

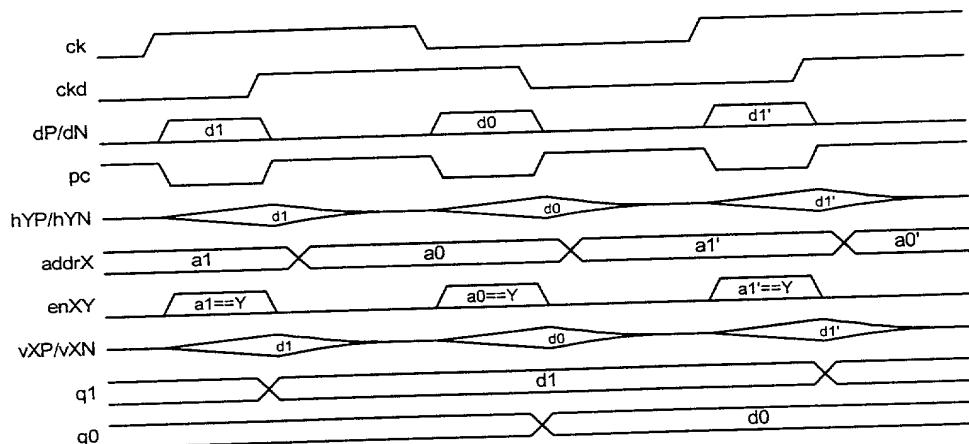


Figure 25